arXiv:1411.4074v1 [math.ST] 14 Nov 2014

Improving Monte Carlo randomized approximation schemes

Mark Huber
Claremont McKenna College
mhuber@cmc.edu

Version: November 18, 2014

Abstract

Consider a central problem in randomized approximation schemes that use a Monte Carlo approach. Given a sequence of independent, identically distributed random variables $X_1, X_2, \ldots$ with mean μ and standard deviation at most $c\mu$, where c is a known constant, and $\epsilon, \delta > 0$, create an estimate $\hat{\mu}$ for μ such that $\mathbb{P}(|\hat{\mu} - \mu| > \epsilon\mu) \leq \delta$. This technique has been used for building randomized approximation schemes for the volume of a convex body, the permanent of a nonnegative matrix, the number of linear extensions of a poset, the partition function of the Ising model and many other problems. Existing methods use (to the leading order) $19.35(c/\epsilon)^2 \ln(\delta^{-1})$ samples. This is the best possible number up to the constant factor, and it is an open question as to what is the best constant possible. This work gives an easy to apply estimate that only uses $6.96(c/\epsilon)^2 \ln(\delta^{-1})$ samples in the leading order.

1 Introduction

The most common form of randomized approximation algorithm works by finding an X whose mean μ matches the true answer given the input, next drawing $X_1, \ldots, X_k$ independently and identically distributed (iid) according to X, and finally creating an estimate $\hat{\mu}$ for μ as a function of these random samples.

Applications of this technique include finding the partition function of a Gibbs distribution [5], approximating the number of linear extensions of a poset [1], estimating the volume of a convex body [10], approximating the permanent of a nonnegative matrix [8], approximating the normalizing con-

stant for the ferromagnetic Ising model [7], finding the maximum likelihood for spatial point process models [6], and many others.

Say that such an estimate $\hat{\mu}$ for μ is an (ϵ, δ)-*randomized approximation scheme* (or (ϵ, δ)-ras) if

$$\mathbb{P}(|\hat{\mu} - \mu| > \epsilon\mu) \leq \delta. \tag{1}$$

That is, the chance that the absolute relative error in the estimate is greater than ϵ is at most δ in an (ϵ, δ)-ras. The goal is to create an (ϵ, δ)-ras using $X_1, \ldots, X_T$, where T is a random variable that has a small mean.

For each of the applications mentioned earlier, it was shown how to build random variables X_i such that $\mathrm{SD}(X_i) \leq c\mu$ for a known constant c that is an easily computable function of the input. Given this restriction on the standard deviation, it is well known how to generate an (ϵ, δ)-ras using at most $19.35c^2\epsilon^{-2}\ln(\delta^{-1})$ samples (plus lower order terms). The details are discussed further in Section 2.

On the other hand, it is known from an application of Wald's sequential ratio test that any such algorithm requires at least $\Omega(c^2\epsilon^{-2}\ln(\delta^{-1}))$ samples on average (as shown in [2]), therefore this is the best possible up to the constant factor. The question remains of what is the best constant factor.

This work introduces a simple new algorithm for estimating μ that reduces this constant from 19.35 to 6.96. For all of the applications listed above, this approach immediately improves the constant of the running time by a factor of 2.78.

First define a nuisance factor that will appear in several results. Let

$$f(\epsilon) = (1-\epsilon)^{-2}(1+\epsilon-\epsilon^2)^{-1}(1+\epsilon). \tag{2}$$

Since $f(\epsilon) = 1 + 2\epsilon + O(\epsilon^2)$, the leading order terms of ϵ^{-2} and $\epsilon^{-2}f(\epsilon)$ are identical.

Theorem 1. *Suppose $X_1, X_2, \ldots$ is an independent, identically distributed sequence of random variables with mean μ and standard deviation at most $c\mu$, where c is a known constant. Then for $\epsilon \in (0, 1/3)$ and $\delta \in (0, 1)$, it is possible to find $\hat{\mu}$ such that $\mathbb{P}(|\hat{\mu} - \mu| > \epsilon\mu) \leq \delta$ using $X_1, \ldots, X_t$, where*

$$t = \left\lceil \left(\frac{c}{\epsilon}\right)^2 f(\epsilon) \right\rceil \left(2 \left\lceil \frac{\ln(2\delta^{-1})}{\ln(4/3)} \right\rceil + 1 \right). \tag{3}$$

Note $2/\ln(4/3) \leq 6.96$, which gives the factor in the leading order term mentioned earlier.

2 Estimating μ

The classic estimate for μ uses sample averages of the X_i. Let $S_k = (X_1 + \cdots X_k)/k$. Then $\mathbb{E}[S_k] = \mu$ and $\mathrm{SD}(S_k) = \mathrm{SD}(X)/\sqrt{k}$. Say that the estimate *fails* if the absolute relative error $|(\hat{\mu}/\mu) - 1|$ is greater than ϵ. Using Chebyshev's inequality gives $\mathbb{P}(|(\hat{\mu}/\mu) - 1| > \epsilon) \le c^2/(\epsilon^2 k)$. This gives a bound on the probability of failure that only goes down polynomially in the number of samples k.

2.1 Median of means

A well known estimate with an exponentially small chance of failure goes back to at least [4]. The central idea is to look at the median of several draws, each of which is the average of some fixed number of draws of the original random variable.

Let $k = \lceil 8(c/\epsilon)^2 \rceil$ draws of X_i. So Then $\mathbb{E}[S_k] = \mu$ and $\mathrm{SD}(S_k) = \epsilon\mu/\sqrt{8}$. Then Chebyshev's inequality gives

$$\mathbb{P}(|S_k - \mu| \ge \epsilon\mu) \le 1/8. \tag{4}$$

The next step is to draw $W_1, \ldots, W_{2k+1}$ iid from the same distribution as S_k. Then it is highly likely that the median of the $\{W_i\}$ values falls inside the region that has 7/8 probability.

To be precise:

Lemma 1. *Suppose $\mathbb{P}(R \le a) \le p$ and $\mathbb{P}(R \ge b) \le p$ for some $a < b$. Then for $R_1, R_2, \ldots \sim R$ iid,*

$$\mathbb{P}(\mathrm{median}\{R_1, \ldots, R_{2k+1}\} \notin [a, b]) \le 4(\pi(k+1))^{-1/2}[4p(1-p)]^k. \tag{5}$$

For $\mathbb{P}(R \in [a, b]) \ge 1 - p$ then

$$\mathbb{P}(\mathrm{median}\{R_1, \ldots, R_{2k+1}\} \notin [a, b]) \le 2(\pi(k+1))^{-1/2}[4p(1-p)]^k. \tag{6}$$

Proof. Suppose that $\mathbb{P}(R \in [a, b]) \ge 1 - p$. If $\mathbb{P}(R \in [a, b]) = 1$ the chance the median is not in $[a, b]$ is 0 and the inequality holds.

Otherwise, let U be uniform over $[0, 1]$, $R_{\text{in } [a,b]}$ be the distribution of R conditioned on $R \in [a, b]$, and $R_{\text{not in } [a,b]}$ be the distribution of R conditioned on $R \notin [a, b]$.

Let $\mathbb{1}$(expression) be the indicator function that is 1 when the expression is true, and 0 when it is false. Then for independent random variables U, $R_{\text{in}[a,b]}$, $R_{\text{not in } [a,b]}$,

$$R \sim R_{\text{in } [a,b]}\mathbb{1}(U \le \mathbb{P}(R \in [a, b])) + R_{\text{not in } [a,b]}\mathbb{1}(U > \mathbb{P}(R \in [a, b])).$$

With this representation, the median of $2k+1$ draws of R will fall into $[a,b]$ if at least $k+1$ of the U_i fall into $[0,\mathbb{P}(R\in[a,b])]$, which in turn occurs when the median of the U_i is in $[0,1-p]$.

The median of the U_i (call it M) is well known to have a beta distribution with density $f_M(x)=x^k(1-x)^k\Gamma(2k+2)/\Gamma(k+1)^2$. Note $\mathbb{P}(M>1-p)=\int_{1-p}^1 f_M(x)\,dx=\int_0^p f_M(x)\,dx=\int_0^p[x(1-x)]^k\Gamma(2k+2)/\Gamma(k+1)^2$.

For $x\in[0,p]$, $x(1-x)\le p^2+(1-2p)x$, so

$$\begin{aligned}\mathbb{P}(M>1-p)&\le\int_0^p\frac{[p^2+(1-2p)x]^k\Gamma(2k+1)}{\Gamma(k)^2}\\&=\frac{\Gamma(2k+2)}{\Gamma(k+1)^2}\cdot\left[\frac{(p-p^2)^k-p^{2k}}{(k+1)(1-2p)}\right]\\&=\frac{(2k+2)!}{(k+1)!(k+1)!}\cdot\frac{k+1}{2k+2}\cdot\left[\frac{p^k[(1-p)^k-p^k]}{1-2p}\right]\end{aligned}$$

The factor $(2k+2)!/[(k+1)!]^2$ (known as a central binomial coefficient) is well known to be at most $2^{2k+2}/\sqrt{\pi(k+2)}$ (see [9]). Simplifying and neglecting the $-p^k$ term then gives the result.

The result where both $\mathbb{P}(R\le a)\le p$ and $\mathbb{P}(R\ge p)$ is similar. □

Applying this lemma to W_i with probability $7/8$ of landing in the desired location gives a chance of error at most $[4(7/8)(1-7/8)]^k=\exp(-\ln(7/16)k)$ for $k\ge 5$. Now the chance of failure is declining exponentially.

Lemma 2. *The preceding procedure gives an (ϵ,δ)-ras for μ that uses at most $\lceil 8(c/\epsilon)^2\rceil(2\lceil\ln(\delta^{-1})/\ln(16/7)\rceil+1)$ samples.*

Proof. An instance of W takes $\lceil 8(c/\epsilon)^2\rceil$ draws from X to produce.

To make $[4(7/8)(1-7/8)]^k\le\delta$, $k\ge\ln(\delta^{-1})/\ln(16/7)$. Since $2k+1$ draws of W are necessary, the result follows. □

The method just described could be done with $W\sim S_{\lceil i(c/\epsilon)^2\rceil}$ for any i. The choice of $i=8$ minimizes the constant in the running time given in the previous lemma. In finding k, the $4(\pi(k+1))^{-1/2}$ factor was bounded by 1 for $k\ge 5$. Using the full factor leaves the first order term unchanged, only affecting lower order terms.

Note $8\cdot(1/\ln(16/7))\cdot 2\approx 19.35$, giving the constant mentioned earlier.

2.2 The new estimate

The new method creates a new random variable V such that

$$\mathbb{P}(|V-\mathbb{E}[V]|\ge\epsilon\mu)\le 1/4, \tag{7}$$

but only using slightly more than $(c/\epsilon)^2$ draws from X_i. To accomplish the same feat using Chebyshev's inequality would require $4(c/\epsilon)^2$ draws from the X_i.

Before describing the procedure, it will help to have an understanding of why a random variable does not always lie inside the standard deviation. Suppose that Y has mean μ and standard deviation $\epsilon\mu$. An example of such a random variable is $\mathbb{P}(Y = \mu - \epsilon\mu) = \mathbb{P}(Y = \mu + \epsilon\mu) = 1/2$.

Note that $\mathbb{P}(Y < \mu + \epsilon\mu) = 1/2$, because fully half of the probability is sitting just outside the interval $(\mu - \epsilon\mu, \mu + \epsilon\mu)$. The goal is to create V such that $\mathbb{P}(V < \mu + \epsilon\mu) \geq 3/4$.

To create such a V, let R be uniform over the interval $[1 - \epsilon, 1 + \epsilon]$, and independent of Y. Then set $V = RY$. So

$$\mathbb{P}(YR < \mu + \epsilon\mu) = \frac{1}{2} + \frac{1}{2}\mathbb{P}(R < 1) = \frac{3}{4} \tag{8}$$

and

$$\mathbb{P}(YR > \mu - \epsilon\mu) = \frac{1}{2} + \frac{1}{2}\mathbb{P}(R > 1) = \frac{3}{4}. \tag{9}$$

So for this particular Y, by applying this simple random scaling procedure, the chance of falling into the tails has a bound as small as when 4 samples from Y are averaged together!

This is the idea behind the new estimate. Given $\epsilon > 0$, for a given draw $S_i \sim S$, draw $R_i \sim \mathsf{Unif}([1 - \epsilon, 1 + \epsilon])$ independently of S_i. Next, let

$$W_i' = S_i R_i. \tag{10}$$

This smoothed random variable still has $\mathbb{E}[W_i'] = \mu$, moreover, it is now much more likely to lie within a standard deviation of its mean!

Lemma 3. *Let S have mean μ and standard deviation at most $\epsilon\mu/\sqrt{f(\epsilon)}$, where $\epsilon \leq 1/3$. Let R be independent of S and uniform over $[1 - \epsilon, 1 + \epsilon]$. Then*

$$\mathbb{P}(SR \leq \mu - \epsilon\mu) \leq \frac{1}{4}, \quad \mathbb{P}(SR \geq \mu + \epsilon\mu) \leq \frac{1}{4}. \tag{11}$$

This lemma is the heart of the new estimate, and will be proved in the next section. Using this lemma together with Lemma 1 immediately gives the following result.

Lemma 4. *Fix $\epsilon \in (0, 1/3)$, t a positive integer, and let X be a random variable with mean μ and standard deviation at most $c\mu$. For $i = 1, 2, \ldots, t$,*

let S_i be the sample average of $\lceil (c/\epsilon)^2 f(\epsilon) \rceil$ iid draws from the distribution of X. Independently, draw $R_1, \ldots, R_t$ uniformly from $[1-\epsilon, 1+\epsilon]$. Let

$$t = \lceil (cf(\epsilon)/\epsilon)^2 \rceil (2\lceil \ln(2\delta^{-1})/\ln(4/3) \rceil + 1). \tag{12}$$

Then it holds that

$$\mathbb{P}(|\operatorname{median}\{S_1 R_1, \ldots, S_t R_t\} - \mu| > \epsilon\mu) \leq \delta. \tag{13}$$

3 Proof of Lemma 3

The proof of Lemma 3 takes the following steps. Let S be a random variable with mean μ and standard deviation at most $\alpha\mu$. The goal is to show that for $R \sim \mathsf{Unif}([1-\epsilon, 1+\epsilon])$, $\mathbb{P}(SR \leq \mu - \epsilon\mu)$ and $\mathbb{P}(SR \geq \mu + \epsilon\mu)$ are both at most $1/4$.

1. Eliminate the μ factor by considering $Y = S/\mu$ so $\mathbb{E}[Y] = 1$ and $\mathrm{SD}(Y) \leq \alpha$. The new goal is to upper bound $q_+ = \mathbb{P}(YR \geq 1+\epsilon)$ and $q_- = \mathbb{P}(YR \leq 1-\epsilon)$.

2. First look at q_+. Write Y as a mixture of random variables Y_1 and Y_2 where there is a value y such that $\mathbb{P}(Y_1 \leq y) = \mathbb{P}(Y_2 > y) = 1$. (Lemma 5.)

3. Show that for a proper choice of y, the value q_+ is maximized when Y_1 and Y_2 are not random, but deterministic functions of ϵ, and when the standard deviation of Y equals the upper bound $\alpha\mu$. (Lemmas 8 and 9.)

4. Show that for $\mathrm{SD}(Y) = \epsilon/\sqrt{f(\epsilon)}$, q_+ is at most $1/4$ given that Y_1 and Y_2 are each concentrated at a single value. (Lemma 11.)

5. The proof for q_- is then accomplished in a similar fashion.

To begin, note that if $Y < 1$, then $YR < 1+\epsilon$ always. When $Y > (1+\epsilon)/(1-\epsilon)$, then $YR > 1+\epsilon$. Then there is the case in the middle, where $Y \in [1, (1+\epsilon)/(1-\epsilon)]$ and it could be true that $YR \leq 1+\epsilon$ if R is small enough. For that reason, define the intervals $I_1 = (-\infty, 1)$, $I_2 = [1, (1+\epsilon)/(1-\epsilon)]$, $I_3 = ((1+\epsilon)/(1-\epsilon), \infty)$.

Write Y as a mixture of two random variables, one of which has support on I_1, and the other has support on $I_2 \cup I_3$.

Lemma 5. *For any random variable Y, there exist Y', C, Y_1, and Y_2 such that $Y' \sim Y$ and*

$$Y' = CY_1 + (1 - C)Y_2, \tag{14}$$

where Y_1 falls into I_1 with probability 1, Y_2 falls into $I_2 \cup I_3$ with probability 1, and C is a Bernoulli random variable with $\mathbb{P}(C = 1) = \mathbb{P}(Y \in I_1)$ and $\mathbb{P}(C = 0) = \mathbb{P}(Y \in I_2 \cup I_3)$.

Proof. If either $\mathbb{P}(Y \in I_1)$ or $\mathbb{P}(Y \in I_2 \cup I_3)$ equal 1, the result is trivial, otherwise, let $Y_1 \sim [Y|Y \in I_1]$, $Y_2 \sim [Y|Y \in I_2 \cup I_3]$, and the result follows. □

Since $Y' \sim Y$, $\mathbb{P}(YR \leq 1 + \epsilon) = \mathbb{P}(Y'R \leq 1 + \epsilon)$. So without loss of generality work with Y' from here on out.

Recall that the probability that a event occurs is just the expected value of the indicator function of that event. So

$$\mathbb{P}(Y'R \leq 1 + \epsilon) = \mathbb{E}[\mathbb{1}(Y'R \leq 1 + \epsilon)]. \tag{15}$$

We also need the following well known fact about conditional expectation (see for instance [3]).

Fact 1. *If A and B are random variables such that A and $[A|B]$ are both integrable, then*

$$\mathbb{E}[A] = \mathbb{E}[\mathbb{E}[A|B]]. \tag{16}$$

Hence

$$\mathbb{P}(Y'R \leq 1 + \epsilon) = \mathbb{E}[\mathbb{E}[\mathbb{1}(Y'R \leq 1 + \epsilon)|Y'] \tag{17}$$

and our analysis can start with the inside expectation $\mathbb{E}[\mathbb{1}(Y'R \leq 1+\epsilon)|Y']$.

Lemma 6.

$$\mathbb{E}[\mathbb{1}(Y'R \leq 1+\epsilon|Y'] = \mathbb{1}(Y' \in I_1) + \mathbb{1}(Y' \in I_2)[((1+\epsilon)/Y') - (1-\epsilon)](2\epsilon)^{-1}. \tag{18}$$

Proof. Let $f(R, Y') = \mathbb{1}(R \leq (1 + \epsilon)/Y')$. Then

$$f(R, Y') = f(R, Y')\mathbb{1}(Y' \in I_1) + f(R, Y')\mathbb{1}(Y' \in I_2) + f(R, Y')\mathbb{1}(Y' \in I_3). \tag{19}$$

As noted earlier, $f(R, Y')\mathbb{1}(Y' \in I_3) = 0$ and $f(R, Y')\mathbb{1}(Y' \in I_1) = \mathbb{1}(Y' \in I_1)$. That means

$$\mathbb{E}[f(R, Y')|Y'] = \mathbb{E}[\mathbb{1}(Y' \in I_1)|Y'] + \mathbb{E}[\mathbb{1}(Y' \in Y_2)f(R, Y')|Y']. \tag{20}$$

Since $\mathbb{1}(Y' \in I_2)$ is measurable with respect to Y', $\mathbb{E}[\mathbb{1}(Y' \in I_1)|Y'] = \mathbb{1}(Y' \in I_1)$. In the expectation $\mathbb{E}[\mathbb{1}(Y' \in I_2)f(R, Y')|Y']$, treat Y' as a constant. Since R is uniform over $[1-\epsilon, 1+\epsilon]$, the chance that it is at most $(1-\epsilon)/Y'$ when $Y' \in I_2$ is $[[(1+\epsilon)/Y'] - (1-\epsilon)]/[(1+\epsilon) - (1-\epsilon)]$. Hence

$$\mathbb{E}[f(R, Y')|Y'] = \mathbb{1}(Y' \in I_1) + \mathbb{1}(Y' \in I_2)[((1+\epsilon)/Y') - (1-\epsilon)](2\epsilon)^{-1}. \quad (21)$$

□

Lemma 7. *Let $p = \mathbb{P}(Y' \in I_1)$. Then*

$$\mathbb{E}[\mathbb{1}(RY' \leq 1+\epsilon)] = p + (1-p)\mathbb{E}\left[\mathbb{1}(Y_2 \in I_2)\frac{((1+\epsilon)/Y_2) - (1-\epsilon)}{2\epsilon}\right]. \quad (22)$$

Proof. From the last lemma

$$\mathbb{E}[f(R, Y')] = \mathbb{E}[\mathbb{1}(Y' \in I_1)] + \mathbb{E}[\mathbb{1}(Y' \in I_2)[((1+\epsilon)/Y') - (1-\epsilon)](2\epsilon)^{-1}].$$

Now $\mathbb{E}[\mathbb{1}(Y' \in I_2)] = \mathbb{P}(Y' \in I_2) = p$. For the second term,

$$\begin{aligned}
&\mathbb{E}[\mathbb{1}(Y' \in I_2)[((1+\epsilon)/Y') - (1-\epsilon)](2\epsilon)^{-1}] \\
&\quad = \mathbb{E}[\mathbb{E}[\mathbb{1}(Y' \in I_2)[((1+\epsilon)/Y') - (1-\epsilon)](2\epsilon)^{-1}|C]] \\
&\quad = \mathbb{P}(C=1)(0) + \mathbb{P}(C=0)\mathbb{E}[\mathbb{1}(Y_2 \in I_2)[((1+\epsilon)/Y_2) - (1-\epsilon)](2\epsilon)^{-1}|C]].
\end{aligned}$$

Using $\mathbb{P}(C=0) = 1-p$ completes the proof. □

Recall Jensen's inequality.

Fact 2 (Jensen's inequality)**.** *If X is a random variable with finite mean, $\mathbb{P}(X \in A) = 1$, and g is a convex measurable function over A, then*

$$\mathbb{E}[g(X)] \geq g(\mathbb{E}[X]).$$

Now $g(x) = \mathbb{1}(x \in I_2)((1+\epsilon)/x - (1-\epsilon)(2\epsilon)^{-1}$ is a convex function over all $x \in I_2 \cup I_3$. Similarly, $g(x) = 1$ is a convex function over all $x \in I_1$. That gives the following.

Lemma 8. *Suppose $Y'' = C\mathbb{E}[Y_1] + (1-C)\mathbb{E}[Y_2]$. Then $\mathbb{E}[Y''] = \mathbb{E}[Y'] = 1$ and $\mathrm{SD}(Y'') \leq \mathrm{SD}(Y')$. Also*

$$\mathbb{P}(Y''R \leq 1+\epsilon) \leq \mathbb{P}(Y'R \leq 1+\epsilon). \quad (23)$$

Proof. Equation (23) follows immediately from Jensen's inequality applied to (22). The statement $\mathbb{E}[Y''] = \mathbb{E}[Y']$ follows from $\mathbb{E}[Y'] = \mathbb{E}[\mathbb{E}[Y'|C]]$. Replacing a component of a mixture by its expectation gives a standard deviation at most that of the original random variable, hence $\mathrm{SD}(Y'') \leq \mathrm{SD}(Y')$. □

So there exist a_1 and a_2 such that $Y'' \in \{1 + a_1, 1 + a_2\}$, where $a_1 \leq 0$ and $a_2 > 0$ (since $\mathbb{E}[Y''] = 1$). Note that $\mathrm{SD}(Y'') = \mathbb{P}(Y'' = a_1)a_1^2 + \mathbb{P}(Y'' = a_2)a_2^2$.

Lemma 9. *Let $\mathcal{Y}$ be the set of random variables with mean 1 and standard deviation at most α. Let $\mathcal{W}$ be the set of random variables W with mean 1, standard deviation equal to α, and W takes on one of two values with probability 1. Then*

$$\min_{Y \in \mathcal{Y}} \mathbb{P}(YR \leq 1 + \epsilon) \geq \min_{W \in \mathcal{W}} \mathbb{P}(WR \leq 1 + \epsilon).$$

Proof. Let $\mathcal{Y}''$ be the set of mean 1, standard deviation at most α random variables that only take on one of two values with probability 1. Then the previous lemma got us to

$$\min_{Y \in \mathcal{Y}} \mathbb{P}(YR \leq 1 + \epsilon) \geq \min_{Y'' \in \mathcal{Y}''} \mathbb{P}(Y''R \leq 1 + \epsilon).$$

Suppose $Y'' \in \mathcal{Y}''$ with $p_1 = \mathbb{P}(Y'' = 1 + a_1)$ and $p_2 = 1 - p_1 = \mathbb{P}(Y'' = 1 + a_2)$, where $a_1 \leq 0$ and $a_2 > 0$ and $p_1, p_2 \geq 0$. With this notation, $\mathrm{SD}(Y'') = p_1 a_1^2 + p_2 a_2^2$, and

$$\mathbb{P}(Y''R \leq 1 + \epsilon) = p_1 + p_2(\mathbb{1}(a_2 \in [0, 2\epsilon/(1-\epsilon)]) \left(\frac{((1+\epsilon)/a_2) - (1-\epsilon)}{2\epsilon} \right).$$

This is a decreasing function of a_2 and independent of a_1.

So construct W by setting $k_1 = \alpha a_1 / \mathrm{SD}(Y'')$. $k_2 = \alpha a_2 / \mathrm{SD}(Y'')$. Since $\alpha \geq \mathrm{SD}(Y'')$, $k_2 \geq a_2$, which means $\mathbb{P}(WR \leq 1+\epsilon) \leq \mathbb{P}(Y''R \leq 1+\epsilon)$. Note $\mathbb{E}[W] = 1$ and $\mathrm{SD}(W) = \epsilon$, proving the result. □

What has been accomplished so far it to show that $\mathbb{P}(YR \leq 1 + \epsilon)$ is at least the optimal objective function value for the optimization problem:

$$\begin{aligned} &\min p_1 + p_2(\mathbb{1}(k_2 \in [0, 2\epsilon/(1-\epsilon)])((1 + \epsilon/k_2) - (1-\epsilon))(2\epsilon)^{-1} \\ &\text{subject to } p_1 + p_2 = 1 \\ &\qquad p_1 k_1 + p_2 k_2 = 0 \\ &\qquad p_1 k_1^2 + p_2 k_2^2 = \alpha \end{aligned}$$

Using the constraints to solve for p_1 and p_2 in terms of k_2 turns the objective function into $\min_{k_2>0} h(k_2)$, where

$$h(k_2) = \frac{k_2^2}{k_2^2+\alpha^2} + \frac{\alpha^2}{k_2^2+\alpha^2} \mathbb{1}\left(k_2 \le \frac{2\epsilon}{1-\epsilon}\right) \frac{(1+\epsilon)/(1+k_2) - (1-\epsilon)}{2\epsilon}. \quad (24)$$

Lemma 10. *For $\alpha = \epsilon\sqrt{[(1-\epsilon)^2(1+\epsilon-\epsilon^2)]/[1+\epsilon]}$ the minimum of $h(k_2)$ for $k_2 > 0$ is at least 3/4.*

Proof. When $k_2 \ge 2\epsilon/(1-\epsilon)$, $h(k_2) = k_2^2/(k_2^2+\alpha^2)$, which is an increasing function of k_2. Therefore the minimum occurs at $k_2 = 2\epsilon/(1-\epsilon)$ so it is only necessary to consider $k_2 \in (0, 2\epsilon/(1-\epsilon)]$.

Assuming this to be true and simplifying $h(k_2)$ gives

$$h(k_2) = 1 + \frac{\alpha^2}{k_2^2+\alpha^2}\left[\left(\frac{1+\epsilon}{1+k_2} - (1-\epsilon)\right)(2\epsilon)^{-1} - 1\right]. \quad (25)$$

The right hand side is at least 3/4 if and only if

$$\alpha^2 \le f_1(k_2) \text{ where } f_1(k_2) = \frac{k_2^2(k_2+1)\epsilon}{k_2(\epsilon+2)-\epsilon}. \quad (26)$$

Now

$$\frac{df_1(k_2)}{dk_2} = \frac{2k_2\epsilon(k_2^2(\epsilon+2) + k_2(1-\epsilon) - \epsilon)}{[k_2(\epsilon+2)-\epsilon]^2}.$$

The denominator and the $2k_2\epsilon$ factor in the numerator is always positive, leaving only the quadratic factor $k_2^2(\epsilon+2) + k_2(1-\epsilon) - \epsilon$ to determine the sign. This factor starts negative at $k_2 = 0$, and is then increasing, so the minimum of $f_1(k_2)$ occurs when the quadratic is zero, which happens at $k_2^*(\epsilon) = [\epsilon - 1 + \sqrt{5\epsilon^2+6\epsilon+1}]/(2(\epsilon+2))$. It is easy to bound this expression for $\epsilon \in (0,1)$:

$$\epsilon - \epsilon^2 \le k_2^*(\epsilon) \le \epsilon. \quad (27)$$

Using the lower bound on k_2^* for the numerator and the upper bound on k_2^* for the denominator of $f_1(k_2)$ gives

$$f_1(k_2) \ge \frac{(\epsilon-\epsilon^2)^2(1+\epsilon-\epsilon^2)\epsilon}{\epsilon(\epsilon+2)-\epsilon} = \epsilon^2\frac{(1-\epsilon)^2(1+\epsilon-\epsilon^2)}{1+\epsilon} = \epsilon^2/f(\epsilon). \quad (28)$$

This is why $f(\epsilon)$ was chosen to be what it is. Therefore, for $\alpha^2 \le \epsilon^2/f(\epsilon)$, $\alpha^2 \le f_1(k_2)$, and $h(k_2) \ge 3/4$. □

This lemma shows that $\mathbb{P}(SR \leq (1+\epsilon)\mu) \leq 3/4$ for all S with mean μ and standard deviation at most $\epsilon\mu/f(\epsilon)$. This same sequence of steps, where we first show that we need only consider random variables that take on two values, and then use the constraints on the variable to reduce it to a one dimensional optimization problem, and then finally obtain a bound on the standard deviation.

The result is a function similar to h of (24). Define

$$h_2(k_1) = \frac{k_1^2}{k_1^2+\alpha^2} + \frac{\alpha^2}{k_1^2+\alpha^2}\mathbb{1}\left(k_1 \geq \frac{-2\epsilon}{1-\epsilon}\right)\frac{1+\epsilon-(1-\epsilon)/(1+k_1)}{2\epsilon}. \quad (29)$$

Then $\mathbb{P}(SR \geq (1-\epsilon)\mu) \geq \min_{k_1 \leq 0} h_2(k_1)$.

Lemma 11. *For $\alpha = \epsilon\sqrt{[(1+\epsilon)^3(1-2\epsilon)]/[1-3\epsilon+2\epsilon^2]}$, the minimum of $h_2(k_1)$ for $k_1 \leq 0$ is at least $3/4$.*

The proof is similar to that of Lemma 10.

For $\epsilon \leq 1/3$, the bound on α from Lemma 10 is stronger than the bound from Lemma 11. This proves Lemma 3.

4 Summary

For any random variable with $\mathbb{E}[S] = \mu$ and $\mathrm{SD}(S) \leq \epsilon/f(\epsilon)$, simply generating independently R uniformly over $[1-\epsilon, 1+\epsilon]$ gives SR with the properties that $\mathbb{E}[SR] = \mu$, $\mathbb{P}(SR - \mu \geq \epsilon\mu) \leq 1/4$, and $\mathbb{P}(SR - \mu \leq \epsilon\mu) \leq 1/4$. Since generating S is usually very costly (as in [5, 1, 10, 8, 7, 6]), generating R incurs very little overhead and then allows the median of independent draws to quickly concentrate the resulting approximation.

References

[1] J. Banks, S. Garrabrant, M. Huber, and A. Perizzolo. Using TPA for approximating the number of linear extensions. arXiv:1010.4981. Submitted, 2010.

[2] P. Dagum, R. Karp, M. Luby, and S. Ross. An optimal algorithm for Monte Carlo estimation. *Siam. J. Comput.*, 29(5):1484–1496, 2000.

[3] R. Durrett. *Probability: Theory and Examples, 4th edition.* Cambridge University Press, 2010.

[4] Martin E. Dyer and Alan M. Frieze. Computing the volume of a convex body: A case where randomness provably helps. In Béla Bollobás, editor, *Proceedings of AMS Symposium on Probabilistic Combinatorics and Its Applications*, volume 44 of *Proceedings of Symposia in Applied Mathematics*, pages 123–170. American Mathematical Society, 1991.

[5] M. Huber. Approximation algorithms for the normalizing constant of Gibbs distributions. *Ann. Appl. Probab.* arXiv:1206.2689. To appear.

[6] M. L. Huber and R. L. Wolpert. Likelihood-based inference for Matérn type-III repulsive point processes. *Adv. Appl. Prob.*, 41(4):958–977, 2009.

[7] M. Jerrum and A. Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM J. Comput.*, 22:1087–1116, 1993.

[8] M.R. Jerrum, A. Sinclair, and E. Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *J. of the ACM*, 51(4):671–697, 2004.

[9] T. Koshy. *Catalan Numbers with applications*. Oxford University Press, 2008.

[10] L. Lovász and S. Vempala. Simulated annealing in convex bodies and an $o^*(n^4)$ volume algorithm. *J. Comput. Syst. Sci*, 72(2):392–417.